

IN THE CLAIMS

Please amend the claims to the following.

- 1 1. (Currently Amended) An apparatus ~~system~~ for maintaining cache coherency ~~in a~~
2 ~~CMP~~ comprising:
3 an integrated circuit including
4 a plurality of processor cores, wherein the plurality of processor cores
5 each include a private cache;
6 a shared cache to be shared by the plurality of processor cores, wherein
7 the shared cache includes logic, in response to receiving a write
8 request referencing a block from a requesting processor core of the
9 plurality of processor cores and the block not being owned, to
10 generate a first message including an invalidation part and a write-
11 acknowledgement part, and wherein at least the invalidate part of
12 the first message when received by a second processor core of the
13 plurality of processor cores is to invalidate the block in the second
14 processor core and at least the write-acknowledgement part the
15 first message, when received by the requesting processor core, is
16 also to act as a write acknowledgement to the requesting processor
17 core; and
18 a ring to connect the plurality of processor cores and the shared cache, the
19 ring to transmit the first message to the requesting processor core
20 and second processor core.
- 1 2. (Canceled)

1 3. (Currently Amended) The apparatus ~~system~~ of claim 1 wherein the shared cache
2 includes one or more banks, wherein the one or more cache banks is responsible
3 for a subset of a physical address space of the system, and wherein the block is
4 associated with a physical address of the physical address space of the system.

1 4. (Currently Amended) The apparatus ~~system~~ of claim 1 wherein the first message
2 includes an InvalidateAndAcknowledge message , and wherein the shared cache
3 is to generate the InvalidateAndAcknowledge message, further in response to the
4 block being present in the shared cache and the second processor core being a
5 custodian for the block.

1 5. (Currently Amended) The apparatus ~~system~~ of claim 1 wherein the first message
2 includes an InvalidateAllAndAcknowledge message, and wherein the shared
3 cache, in response to receiving the write request referencing the block from the
4 requesting processor core of the plurality of processor cores and the block not
5 being owned, is to generate the InvalidateAllAndAcknowledge message, further
6 in response to the block not being present in the shared cache and none of the
7 plurality of processor cores being a custodian for the block.

1 6. (Currently Amended) The apparatus ~~system~~ of claim 1 wherein the plurality of
2 processor cores writes data through to the shared cache.

1 7. (Currently Amended) The apparatus ~~system~~ of claim 1 wherein the plurality of
2 processor cores each include a merge buffer, and wherein each of the merge
3 buffers are to coalesce multiple stores to a same block.

- 1 8. (Currently Amended) The apparatus ~~system~~ of claim 1 wherein the shared cache
2 is to fetch a second block from a memory and generate a write acknowledge
3 message to provide a write acknowledgement to the requesting processor core in
4 response to receiving a second write request referencing the second block, the
5 second block not being present in the shared cache and not being owned by any of
6 the plurality of processor cores.
- 1 9. (Currently Amended) The apparatus ~~system~~ of claim 8 wherein the shared cache
2 is to generate an evict message to evict a third block from an owning processor
3 core and generate a second write acknowledge message to provide a second write
4 acknowledgment to the requesting processor core in response to receiving a third
5 write request referencing the third block, the third block being present in the
6 shared cache and the owning processor core of the plurality of cores owns the
7 third block.
- 1 10. (Currently Amended) The apparatus ~~system~~ of claim 1 wherein a bank of the
2 shared cache is to be a home location for a non-overlapping portion of a physical
3 address space associated with the block.
- 1 11. (Currently Amended) The apparatus ~~system~~ of claim 7 wherein each private
2 cache of the plurality of cores are not to hold dirty data, and wherein each of the
3 merger buffers are to hold the dirty data.
- 1 12. (Original) The apparatus ~~system~~ of claim 1 wherein the ring is a synchronous,
2 unbuffered bidirectional ring interconnect.

1 13. (Currently Amended) The apparatus ~~system~~ of claim 12 wherein the first message
2 has a fixed deterministic latency around the ring interconnect.

1 14. (Currently Amended) An apparatus comprising:
2 an integrated circuit including: a plurality of cores and a shared memory
3 connected in a ring, the shared memory to be accessible by each of the
4 plurality of cores, wherein each of the plurality of cores includes a private
5 memory and a merge buffer to purge data to the shared memory, and
6 wherein the shared memory includes receiving logic to receive, from a
7 requesting core of the plurality of cores, a read request referencing the
8 address, ownership logic to determine an owning processor core of the
9 plurality of processor cores owns a block associated with the address, and
10 eviction logic coupled to the receiving logic and the ownership logic, the
11 eviction logic to generate an evict message referencing the an address and
12 associated with the an-owning processor core of the plurality of cores in
13 response to the receiving logic receiving the a-read request refereneing the
14 address from a requesting core of the plurality of cores and the ownership
15 logic determining the owning processor core owns[[ing] the a-block
16 associated with the address.

1 15. (Previously Amended) The apparatus of claim 14, wherein the ring includes a
2 synchronous unbuffered bi-directional ring interconnect.

1 16. (Previously Added) The apparatus of claim 14, wherein the shared memory is a
2 shared cache including a plurality of blocks, and wherein the shared cache is
3 capable of holding each of the plurality of blocks in a cache coherency state.

1 17. (Previously Added) The apparatus of claim 16, wherein the cache coherency
2 state for each of the plurality of blocks is selected from a group consisting of (1) a
3 not present state, (2) a present and owned by a core of the plurality of cores state,
4 (3) a present, not owned, and custodian is a core of the plurality of cores state,
5 and (4) a present, not owned, and no custodian state.

1 18. (Currently Amended) A system comprising:
2 a processor including: a plurality of cores and a shared memory to be coupled
3 together with an unbuffered bi-directional ring interconnect, wherein each
4 of the plurality of cores is to be associated with a private cache memory,
5 the shared memory is to be accessible by each of the plurality of cores,
6 and the shared memory is to include a plurality of blocks, each of the
7 plurality of blocks capable of being held by logic in the shared memory in
8 a not present state; [[,]] a present and owned by a core of the plurality of
9 cores state; [[,]] a present, not owned, and a core of the plurality of cores is
10 a custodian state; [[,]] and a present, not owned, and no core of the
11 plurality of cores is a custodian state; and
12 a system memory associated with the processor to hold elements to be stored by
13 the shared memory.

1 19. (Previously Amended) The system of claim 18, wherein each of the plurality of
2 blocks is a home location for a subset of a physical address space.

1 20. (Previously Amended) The system of claim 19, wherein the shared cache is to
2 generate a first message to invalidate a requested block in all cores of the plurality
3 of cores except for a requesting core of the plurality of cores, in response to
4 receiving a write request referencing the requested block from the requesting core
5 and requested block being held in the present, not owned, and no core of the
6 plurality of cores is a custodian state.

1 21. (New) A method for maintain cache coherency comprising:
2 receiving, with a shared cache, a write request referencing a block from a
3 requesting processor core of the plurality of processor cores on a
4 processor, wherein the plurality of processor cores each include a private
5 cache, and wherein the plurality of cores and the shared cache are
6 connected by a ring interconnect;
7 generating a single message, with the shared cache, in response to receiving the
8 write request;
9 transmitting the single message on the ring interconnect to at least a second
10 processor core of the plurality of processor cores and to the requesting
11 processor core;
12 invalidating the block in the private cache included in the second processor core
13 in response to the second processor core receiving the single message
14 transmitted on the ring interconnect; and
15 write-acknowledging the write request for the requesting processor core in
16 response to the requesting processor core receiving the single message
17 transmitted on the ring interconnect.

1 22. (New) The method of claim 21, wherein the shared cache includes one or more
2 banks, wherein the one or more cache banks is responsible for a subset of a
3 physical address space of a computer system including the processor, and wherein
4 the block is associated with a physical address of the physical address space of the
5 computer system.

1 23. (New) The method of claim 21 wherein the first message includes an
2 InvalidateAndAcknowledge message , and wherein generating the
3 InvalidateAndAcknowledge message, with the shared cache, is further in response
4 to the block being present in the shared cache and the second processor core being
5 a custodian for the block.

1 24. (New) The method of claim 21 wherein the first message includes an
2 InvalidateAllAndAcknowledge message, and wherein generating the
3 InvalidateAllAndAcknowledge message, with the shared cache, is further in
4 response to the block not being present in the shared cache and none of the
5 plurality of processor cores being a custodian for the block.

1 25. (New) The method of claim 21 wherein the plurality of processor cores writes
2 data through to the shared cache.

1 26. (New) The method of claim 21 wherein the plurality of processor cores each
2 include a merge buffer, and wherein each of the merge buffers are to coalesce
3 multiple stores to a same block.

1 27. (New) The method of claim 21, further comprising fetching, with the shared
2 memory, a second block from a memory and generating, with the shared memory,

3 a write acknowledge message to provide a write acknowledgement to the
4 requesting processor core in response to receiving a second write request
5 referencing the second block, the second block not being present in the shared
6 cache and not being owned by any of the plurality of processor cores.

1 28. (New) The method of claim 27 further comprising generating, with the shared
2 cache, an evict message to evict a third block from an owning processor core of
3 the plurality of processor cores and generating a second write acknowledge
4 message to provide a second write acknowledgment to the requesting processor
5 core in response to receiving a third write request referencing the third block, the
6 third block being present in the shared cache and the owning processor core of the
7 plurality of cores owns the third block.

1 29. (New) The method of claim 21 wherein a bank of the shared cache is to be a
2 home location for a non-overlapping portion of a physical address space
3 associated with the block.

1 30. (New) The method of claim 26 wherein each private cache including in the
2 plurality of cores are not to hold dirty data, and wherein each of the merger
3 buffers are to hold the dirty data.

1 31. (New) The method of claim 21 wherein the ring interconnect includes a
2 synchronous, unbuffered, bidirectional, ring interconnect.

1 32. (New) The method of claim 21 wherein the first message has a fixed
2 deterministic latency around the ring interconnect.